

Maximizing the area of intersection of rectangles

D.B.A. Epstein* M.S. Paterson†

24 March 2017

Abstract

This paper attacks the following problem. We are given a large number N of rectangles in the plane, each with horizontal and vertical sides, and also a number $r < N$. The given list of N rectangles may contain duplicates. The problem is to find r of these rectangles, such that, if they are discarded, then the intersection of the remaining $(N - r)$ rectangles has an intersection with as large an area as possible. We will find an upper bound, depending only on N and r , and not on the particular data presented, for the number of steps needed to run the algorithm on (a mathematical model of) a computer. In fact our algorithm is able to determine, for each $s \leq r$, s rectangles from the given list of N rectangles, such that the remaining $(N - s)$ rectangles have as large an area as possible, and this takes hardly any more time than taking care only of the case $s = r$. Our algorithm extends to d -rectangles—analogs of rectangles, but in dimension d instead of in dimension 2. Our method is to exhaustively examine all possible intersections—this is much faster than it sounds, because we do not need to examine all $\binom{N}{s}$ subsets in order to find all possible intersection rectangles. For an extreme example, suppose the rectangles are nested, e.g., concentric squares of distinct sizes, then the only intersections examined are the smallest $s + 1$ rectangles.

1 Background

This problem arose from the use of a novel robotically controlled microscope technology (*Toponome Imaging System*, abbreviated to *TIS*), invented by Walter Schubert (see [4]). TIS records, pixel by pixel, the location and abundance of proteins in a tissue section, thus co-locating many different proteins in the same pixel. Since proximity is necessary (though not sufficient) for interaction, this gives a powerful new method for discovering protein complexes and protein networks, particularly since the anatomy of the section is not disrupted by the TIS process. A TIS run results in a large number (typically, hundreds) of different images of the same tissue section.

Such sequences of images may issue from the TIS machine approximately aligned, but they will seldom be completely correctly aligned without further processing, because physical reality, for example changes in temperature or vibrations from a passing heavy goods vehicle, ensures that camera measurements are never perfect. In order to obtain good information about co-location of proteins and other biomolecules, alignments that

*Mathematics Institute, University of Warwick. Partial support from BBSRC Grant BB/K018868/1.

†Computer Science Department, University of Warwick. Partial support from Centre for Discrete Mathematics and its Applications (DIMAP).




Figure 1: $N = 6$, $r = 2$. Shading shows maximal intersections for $s = 0, 1, 2$.

are as accurate as possible must first be achieved. In [3], it is explained how this can be done. The only adjustments necessary to achieve good alignments turn out to be rigid translation of the different images against one fixed image.

After alignment, some pixels will appear in all N images, and others may appear in only $(N - 1)$ images, or in $(N - 2)$ images, and so on, as shown in Figure 1. If we insist that a pixel can be analyzed only if it records all these signals simultaneously, then we have to restrict to the region where all N images overlap. However, if one is prepared to discard completely one or more images, then the set of pixels which figure in all remaining images will be larger, much larger if one is discarding very seriously misaligned images. In practice one would normally decide to jettison poorest quality images first, so the algorithm presented here, though useful, is not the final word.

For a similar scenario in more conventional photography, suppose one wants an image of St. Mark's Basilica in Venice, but without the tourists and pigeons. To achieve this, one might take snapshots periodically from the same spot over an extended period. The tourists and pigeons can then be eliminated by aligning the images and then averaging the results over each pixel in the intersection of the images. Using our algorithm one can quickly determine which images to discard, in order to obtain a final image of desired area.

2 Statement of result

We start by fixing our notation and clarifying the meaning of the terms we use. In general, we will work in dimension d . In the discussion above, it was assumed that $d = 2$, but our algorithm will work well in general.

Definition 2.1. Given an ordered pair (λ, ρ) of real numbers, we can form the space $\{x : \lambda \leq x \leq \rho\} \subseteq \mathbb{R}$. If $\lambda < \rho$, this is a closed interval; if $\lambda = \rho$, it is a point; and if $\lambda > \rho$ it is the empty set. For the purposes of this paper, it is convenient to have a concept that includes the *endpoints* λ and ρ as well as the underlying space. Refraining from pedantic rigour, we will use the word *interval* as though it is nothing more the underlying space, but nevertheless, we will feel free to extract the “endpoints” whenever necessary, even when $\lambda > \rho$. We will denote such an “interval” by $[\lambda, \rho]$, even if $\lambda > \rho$. \square

Definition 2.2. By a d -rectangle, we mean a product

$$[\lambda_1, \rho_1] \times \cdots \times [\lambda_d, \rho_d] \subseteq \mathbb{R}^d.$$

With the same level of informality as in Definition 2.1, we are able to extract the $2d$ endpoints, even when the underlying space is the empty set. Given a d -rectangle R as above, suppose that $\lambda_i < \rho_i$ for exactly k values of i , and that $\lambda_i = \rho_i$ for the other values of i . Then we call the R a k -dimensional rectangle and define $\dim(R) = k$. \square

Note that a d -rectangle may or may not be a d -dimensional rectangle. In general the underlying space of a d -rectangle may be the empty set or a point or a k -dimensional rectangle in \mathbb{R}^d . A d -rectangle can be presented as a $2d$ -dimensional row vector

$$R = [\lambda_1 \ \rho_1 \ \dots \ \lambda_d \ \rho_d]. \quad (1)$$

Definition 2.3. Given a d -rectangle R as above, we extend the definition of its volume slightly. The d -volume of R is a pair $(\dim(R), \text{vol}(R))$ where:

1. If, for each i , $\lambda_i < \rho_i$, then

$$\dim(R) = d \text{ and } \text{vol}(R) = (\rho_1 - \lambda_1) \times \dots \times (\rho_d - \lambda_d) > 0.$$

2. If $d > k > 0$ and R is a k -dimensional rectangle in \mathbb{R}^d (as in Definition 2.2), then $\dim(R) = k$ and $\text{vol}(R)$ is equal to the volume of R , regarded as a k -rectangle in \mathbb{R}^k . For example, if $k = 1$, then R is an interval in \mathbb{R}^d and $\text{vol}(R)$ is its usual length.

3. If R is a 0-rectangle (that is, a point), then $(\dim(R), \text{vol}(R)) = (0, 0)$.

4. If, for some i , $\lambda_i > \rho_i$, so that the underlying point set of R is \emptyset , then

$$(\dim(R), \text{vol}(R)) = (-1, 0).$$

If $d = 1$, this gives the length of the interval, and if $d = 2$ the area of the rectangle. If d -volumes need to be ordered, then we use lexicographical ordering on the pair $(\dim(R), \text{vol}(R))$. \square

Convention 2.4. Throughout this paper (unless otherwise stated) $d > 0$ will denote the dimension of the euclidean space \mathbb{R}^d in which we are working, and \mathcal{R} will denote a specific ordered list of N d -rectangles. This list can equally well be represented by an $N \times 2d$ real matrix. We fix an integer r with $0 < r < N$.

Our result can now be stated. Its proof will occupy most of the rest of this paper.

Theorem 2.5. *We are given data as in Convention 2.4. Then there is an algorithm taking at most*

$$O\left(d \binom{r+2d}{2d} + dN\right)$$

steps, that determines, for each s with $0 < s \leq r$, exactly s entries to delete from the list \mathcal{R} , so that the list of the remaining $(N - s)$ d -rectangles, represented by the remaining $(N - s)$ elements of \mathcal{R} , has an intersection whose d -volume is maximal, as we vary over the $\binom{N}{s}$ possible choices of s elements of \mathcal{R} . The constant in our bound for the number of steps is independent of N , r and d .

Corollary 2.6. *For any fixed dimension d , the running time of the algorithm is $O(r^{2d} + N)$.*

In applications such as that of microscopy, which was our primary motivation, many rectangles may coincide in some coordinates. Our algorithm is designed to deal efficiently with this situation. Several of the complications result from this.

We have been thinking mainly of the case when r is small compared with N . For r near N , it is possible to achieve better results with a brute force search. For example, if $r = N - 1$, then one deletes all except one d -rectangle. So one need only compute the volume of each of the given d -rectangles, and one then takes the maximum of these. This takes $O(dN)$ steps. We will not bother to make explicit other similar brute force results when r is nearly as big as N .

To be explicit, we need a model of computation: we will assume the usual model of a RAM machine. This means that the number of memory cells is as large as we need for the computation, and that each of these is accessed in unit time. Moreover, memory cells are assumed to be large enough so that any particular integer that needs to be stored can be stored in a single cell. Formally, we do not work with floating point numbers. If necessary, these can be converted to integers by a suitable change of scale. However, it is convenient and less confusing to pretend that we can work with arbitrary real numbers. Arithmetical operations including simple comparisons are assumed to take unit time. This gives a reasonably realistic mathematical model of a computer for most practical purposes.

3 Some notations and definitions

We can think of a d -rectangle as an element in \mathbb{R}^{2d} , a row vector as in Equation 1. \mathcal{R} is then a map $\mathcal{R} : \{1, \dots, N\} \rightarrow \mathbb{R}^{2d}$. In particular, repetition is allowed.

Convention 3.1. Formally speaking, a list is a map, like \mathcal{R} . However, it is convenient to use the same notation also for the image of the map. The advantage of this abuse of notation is that it allows us to write $R \in \mathcal{R}$ for any element R in the list. This paper will contain several lists, and we allow the same abuse of notation in each case. We will disambiguate if we think there is a possibility of confusion.

For $1 \leq j \leq N$, we write R_j instead of $\mathcal{R}(j)$, and, following Equation 1, we will consistently use the notation:.

$$R_j = [\lambda_{j,1} \quad \rho_{j,1} \dots \quad \lambda_{j,d} \quad \rho_{j,d}] = [r_{j,1} \quad \dots \quad r_{j,2d}] \quad (2)$$

Definition 3.2. We often refer to elements of $\{1, \dots, N\}$ as *rectangle indices*, to remind ourselves that they are indices for our list \mathcal{R} of d -rectangles. \square

Definition 3.3. Let $L \subseteq \{1, \dots, N\}$ be non-empty. The *intersection* of the family $(R_j)_{j \in L}$ is defined as follows. For each i with $1 \leq i \leq d$, let $M_i = \max_{j \in L} \lambda_{ji}$ and let $m_i = \min_{j \in L} \rho_{ji}$. Then the *intersection* is the d -rectangle

$$R(L) = [M_1 \quad m_1 \quad \dots \quad M_d \quad m_d]. \quad (3)$$

Equation 3 is the result of a trivial calculation when $R(L)$ is not empty, and is a definition when the point set of $R(L)$ is empty. $R(L)$ is called an *intersection rectangle*. We denote by \mathcal{J}_r the set of non-empty intersection rectangles ($R(L) \neq \emptyset$), where $L \subseteq \{1, \dots, N\}$ and $|L| \geq N - r$. \square

We will find an efficient way to produce, one by one, the elements of \mathcal{J}_r . We will then be able to calculate, for each $R \in \mathcal{J}_r$, the d -volume of R and compare these results to find the best one. The number of subsets like L above, giving rise to elements of \mathcal{J}_r in Definition 3.3, could be as large as $\sum_{s=0}^r \binom{N}{s}$, potentially a large number. To reduce this number, note that a given element $R \in \mathcal{J}_r$ can in general satisfy $R = R(L)$ for many different $L \subset \{1, \dots, N\}$. For any such L , we define $L(R)$, by

$$L(R) = \{j | R \subseteq R_j\}, \quad (4)$$

which is the largest subset of the given set of rectangles with intersection R . Then $L \subseteq L(R) = L(R(L))$, and $|L(R(L))| \geq |L| \geq N - r$.

We are thinking of r as small compared with N , and so it is easier to work with the complement of $L(R)$ in $\{1, \dots, N\}$. We define

$$U(R) = \{j | R \not\subseteq R_j\} \quad (5)$$

so that $|U(R)| \leq r$.

Definition 3.4. We define

$$\mathcal{U}_r = \{U(R) : R \in \mathcal{J}_r\}.$$

The map $U : \mathcal{J}_r \rightarrow \mathcal{U}_r$ is clearly a bijection. \square

An essential feature of our solution will be the use of different orderings on the set $\{1, \dots, N\}$, one for each i with $1 \leq i \leq 2d$. To understand better why and how orderings are important, we discuss the special case $d = 1$. In this case, we are given a list of N intervals $\mathcal{R} = (I_1, \dots, I_N)$, with $I_j = [\lambda_j, \rho_j]$. To avoid irrelevant detail, we assume for the moment that the $2N$ endpoints of these intervals are all distinct, and that the intersection of the N intervals is not empty. We define $u, v \in \{1, \dots, N\}$ so that $\lambda_u = \max_j \lambda_j = m$ and $\rho_v = \min_j \rho_j = M$, as in Definition 3.3. Unless I_u is removed from \mathcal{R} , the lefthand endpoint of the intersection will continue to be m . Unless I_v is removed from \mathcal{R} , the righthand endpoint of the intersection will continue to be M . So I_u or I_v must be removed (and remember that $u = v$ is a possibility) if the length of the intersection is to strictly increase. If $r \geq 2$, and we are omitting r intervals from \mathcal{R} , then similar reasoning will apply to subsequent removal of intervals. We have to look repeatedly for the largest lefthand endpoint and, separately, for the smallest righthand endpoint, in the remaining intervals.

Given a list \mathcal{R} of N d -rectangles, we now define, for each i with $1 \leq i \leq 2d$, orderings \prec_i and \preceq_i on $\{1, \dots, N\}$, modelling the definitions on the discussion in the previous paragraph.

Definition 3.5. Let $j, k \in \{1, \dots, N\}$ and let $1 \leq t \leq d$. If $i = 2t - 1$, we define $j \prec_i k$ if $\lambda_{j,t} > \lambda_{k,t}$, and $j \preceq_i k$ if $\lambda_{j,t} > \lambda_{k,t}$ or if $(\lambda_{j,t} = \lambda_{k,t} \text{ and } j \leq k)$. If $i = 2t$, we define $j \prec_i k$ if $\rho_{j,t} < \rho_{k,t}$, and $j \preceq_i k$ if $\rho_{j,t} < \rho_{k,t}$ or if $(\rho_{j,t} = \rho_{k,t} \text{ and } j \leq k)$. Then \prec_i is a partial order and \preceq_i is a total order on $\{1, \dots, N\}$. For odd i , the orderings are (weakly) decreasing in the λ 's, and, for even i , they are (weakly) increasing in the ρ 's. \square

4 Properties of $X \in \mathcal{U}_r$.

We want to produce the elements $X \in \mathcal{U}_r$ efficiently one by one. With this in mind, let $X \subset \{1, \dots, N\}$ with $|X| \leq r < N$, in which case we find necessary and sufficient conditions for X to be an element of \mathcal{U}_r .

Definition 4.1. An *endpoint list* is a non-empty subset of $\{1, \dots, N\}$, arranged in order, smallest first, with respect to one of the total orderings \preceq_i described in Definition 3.5. We set E_i equal to the set $\{1, \dots, N\}$, arranged in order according to \preceq_i , so that E_i is an endpoint list. \square

Definition 4.2. Let $X \subset \{1, \dots, N\}$, with $|X| \leq r < N$. For each i , let $j(i, X)$ be the smallest element of E_i that is not in X (which must exist since X has fewer than N elements). We define $\text{bv}(i, X) = r_{j(i, X), i}$ (where bv stands for *barrier value*). We set

$$\text{bv}(X) = (\text{bv}(1, X), \dots, \text{bv}(2d, X)). \quad (6)$$

\square

Lemma 4.3. Suppose $X \subset \{1, \dots, N\}$, with $|X| \leq r < N$. $X \in \mathcal{U}_r$ if and only if the following two conditions are satisfied:

4.3.1: For each i with $1 \leq i \leq d$, $\text{bv}(2i-1, X) \leq \text{bv}(2i, X)$.

4.3.2: For each $j \in X$, there is an i such that $j \prec_i j(i, X)$.

Under these conditions, $X = U(\text{bv}(X))$.

Proof. Suppose first that $X \in \mathcal{U}_r$, or equivalently that $X = U(R)$ for some $R \in \mathcal{I}_r$, where

$$\emptyset \neq R = [r_1 \ r_2 \ \dots \ r_{2d-1} \ r_{2d}].$$

For any $j \notin U(R)$, we have $R \subseteq R_j$, and so, for $1 \leq i \leq d$,

$$r_{j, 2i-1} \leq r_{2i-1} \leq r_{2i} \leq r_{j, 2i}.$$

This holds in particular with $j = j(2i-1, X)$ and with $j = j(2i, X)$, giving rise to

$$\text{bv}(2i-1, X) \leq r_{2i-1} \leq r_{2i} \leq \text{bv}(2i, X), \quad (7)$$

proving the first condition in the statement of Lemma 4.3.

Continuing with the hypotheses $R \in \mathcal{I}_r$ and $X = U(R)$, we next show that $\text{bv}(X) = R$, when we regard both as elements of \mathbb{R}^{2d} . Suppose $1 \leq i \leq 2d$. Since $R = \bigcap \{R_j : R \subseteq R_j\}$, we know that $r_i = r_{j, i}$, for some j such that $R \subseteq R_j$, that is, for some $j \notin X$. By Equation 7, if i is odd, we have

$$r_i = r_{j, i} \geq \text{bv}(i, X) = r_{j(i, X), i}.$$

But $j(i, X)$ is, by definition, the smallest element not in X , and the ordering is reversed for odd i , according to the ordering \preceq_i (see Definition 3.5). So a strict inequality $\text{bv}(i, X) < r_{j, i}$ is impossible, and we must have the equality $\text{bv}(i, X) = r_i$. Similarly if i is even. This proves that $R = \text{bv}(X)$. Therefore $X = U(R) = U(\text{bv}(X))$.

We continue to assume that $X = U(R)$, and let $j \in X$. Then $R \not\subseteq R_j$. It follows that, for some i , $[r_{2i-1}, r_{2i}] \not\subseteq [r_{j, 2i-1}, r_{j, 2i}]$. This implies that $r_{2i-1} < r_{j, 2i-1}$ or $r_{2i} > r_{j, 2i}$. We have shown that $R = \text{bv}(X)$, so this means that $j \prec_{2i-1} j(2i-1, X)$ or $j \prec_{2i} j(2i, X)$, proving the second condition in the statement of Lemma 4.3.

Now let us assume the two conditions in the statement of the lemma, and prove that $X \in \mathcal{U}_r$. Let $R = \text{bv}(X)$ be a d -rectangle. By Condition 4.3.1, $R \neq \emptyset$. If $j \notin X$, then, for each i , $j(i, X) \preceq_i j$. Therefore $\text{bv}(2i-1, X) \geq r_{j, 2i-1}$ and $\text{bv}(2i, X) \leq r_{j, 2i}$. It follows that $R \subseteq R_j$. In particular, for each $j \notin X$ and each i with $1 \leq i \leq 2d$, $R \subseteq R_{j(i, X)}$, so that $R \subseteq \bigcap_i R_{j(i, X)}$. On the other hand, Definition 3.3 shows that $\bigcap_i R_{j(i, X)} \subseteq R$. We deduce that $R = \bigcap_{j \notin X} R_j$, which shows that $R \in \mathcal{I}_r$.

If, on the other hand, $j \in X$, then, by hypothesis, for some i , $j \prec_i j(i, X)$. Then $r_{j, i} > \text{bv}(i, X)$ if i is odd, and $r_{j, i} < \text{bv}(i, X)$ if i is even. It follows that $R \not\subseteq R_j$. Therefore $X = \{j : R \not\subseteq R_j\}$ and $X \in \mathcal{U}_r$. \square

5 Program paths

Before describing the algorithm that produces all the elements of \mathcal{U}_r one by one, we describe how to generate one particular $U \in \mathcal{U}_r$. From Lemma 4.3, $U = U(\text{bv}(U))$.

For $1 \leq i \leq 2d$, we write $U_i = \{j | j \prec_i j(i, U)\}$, so that, by Lemma 4.3, $U = \bigcup_i U_i$. For $1 \leq i \leq 2d$ let

$$U^i = U_i \setminus \bigcup_{1 \leq k < i} U_k. \quad (8)$$

U is equal to the disjoint union of the U^i . We generate first the elements of U^1 , then the elements of U^2 , and so on.

We prefer to think of these rectangle indices in U as being *discarded* or deleted, rather than generated, since we are thinking of generating the complement of U , $L = \{1, \dots, N\} \setminus U$, and then taking the intersection of the d -rectangles in L . We start with $\{1, \dots, N\}$, discard the rectangle indices in U^1 , then the rectangle indices in U^2 , and so on. To completely specify the order in which the rectangle indices in U are discarded, we insist that the elements of U^i are discarded in increasing \prec_i order, that is, least first.

Definition 5.1. By a Δ -move we mean one of the deletions just mentioned. While working with U^i , Δ deletes the \prec_i -smallest element of E_i remaining. Until all elements of U^i are deleted, this smallest element is an element of U^i . While carrying out such deletions, we say that the algorithm is *focussed at i* . An S -move moves from the *focus index i* to $(i+1)$. In our process, there are only these two types of move, and the process starts with focus at $i = 1$. Δ and S are mnemonic for “delete” and “shift index by one”. At the end of the process, a final S -move causes the process to stop. \square

During the algorithm the E_i keep on changing. In order to keep track of what is going on, we introduce a time variable t which increases by one at each move. Let $\delta(t, U)$ be the number of Δ -moves and $s(t, U)$ the number of S -moves during the first t units of time, so that $t = \delta(t, U) + s(t, U)$. At time t , the focus is at $i = s(t, U) + 1$. Recall from Definition 4.1 that E_i consists of the elements of $\{1, \dots, N\}$ listed in ascending order according to \prec_i . We write $E_i(t, U)$ for E_i at time t , when all the lists $E_i(t, U)$ have exactly the same set of $N - \delta(t, U)$ entries, though the orders in which these entries appear are, in general, different. When a Δ -move deletes an element of $E_i(t, U)$ during focus at i (always the smallest element according to \prec_i), then that element is also deleted from all the $E_k(t, U)$, for $1 \leq k \leq 2d$.

The process starts at $t = 0$ with focus at $i = 1$. For each k with $1 \leq k \leq 2d$, $E_k(0, U) = \{1, \dots, N\}$. For $1 \leq i \leq 2d$, let t_i be the first time at which the focus moves to i (so $t_1 = 0$), and let t_{2d+1} be the time when the process ends. At time t_i , the set of rectangle indices so far deleted from $\{1, \dots, N\}$ and from the various E_ℓ , is equal to $\bigcup_{1 \leq k < i} U^k$. From the definition of U_i immediately preceding Equation 8, we know that

$$U^i = \{j | j \in E_i(t_i, U) \text{ and } j \prec_i j(i, U)\} \quad (9)$$

(see Convention 3.1 for the meaning of $j \in E_i$), so that the \prec_i -smallest element of U^i , which is the next rectangle number to be discarded (assuming U^i is not empty), is also the \prec_i -smallest element of $E_i(t, U)$. Therefore, $|U^i|$ successive Δ -moves will delete exactly the elements of U^i , and, following this, an S -move transfers focus to $(i+1)$, or, if $i = 2d$, S causes the process to stop.

Rules 5.2. It is convenient to think of the process $\pi(U)$ just described, when $U \in \mathcal{U}_r$, as a word $W(U)$ in the terms S and Δ , read from left to right. Let $u(i) = |U^i|$ be the number of elements in U^i . Note that $\sum_i u(i) \leq r$. We define

$$W(U) = \Delta^{u(1)} S \Delta^{u(2)} S \dots S \Delta^{u(2d)} S. \quad (10)$$

We will now present several rules on words w in S and Δ that every $w = W(U)$ has to obey. Once the rules have all been made explicit, we will define \mathcal{W}_r to be the set of legal words, that is, the set of all words satisfying all the rules. We will then show that the map $F_{WU} : \mathcal{U}_r \rightarrow \mathcal{W}_r$ given by $F_{WU}(U) = W(U)$ is a bijection, so that the rules will be not only necessary but also sufficient for a word w to be equal to $W(U)$ for some $U \in \mathcal{U}_r$. This will reduce our task to producing an algorithm that produces one by one the set of all legal words.

We already have three rules, satisfied by any word $w = W(U)$ for some $U \in \mathcal{U}_r$:

Rule A: w contains exactly $2d$ terms equal to S .

Rule B: w ends with an S .

Rule C: w contains at most r terms equal to Δ .

Rules A and B mean that each coordinate direction is considered and Rule C ensures that at most r deletions are used.

Definition 5.3. Now suppose a word w in Δ and S satisfies Rule A, Rule B and Rule C. We obtain a process $\pi(w)$, analogous to the process $\pi(U)$ (see first sentence of 5.2), that reads w from left to right, taking action as appropriate for each term read. More precisely, we proceed as follows. If $t \leq \text{length}(w)$, we factorize $w = w_t.v_t$, where the length of w_t is t . We denote by $\delta(t, w)$ the number of terms equal to Δ in w_t , and by $s(t, w)$ the number of terms equal to S in w_t . After reading w_t , the focus is at $i = i(t, w) = s(t, w) + 1$. In particular, the focus is at $i = 1$ when $t = 0$. We denote by $E_k(t, w)$ the list E_k after the $\delta(t, w)$ deletions that occur while w_t is being read. With the focus at i , let j be the \preccurlyeq_i -smallest element of $E_i(t, w)$. If the first term of v_t is Δ , then j is deleted from each $E_k(t, w)$, giving the endpoint list $E_k(t+1, w)$. If the first term of v_t is S , then focus is transferred from i to $(i+1)$, and, for each k with $1 \leq k \leq 2d$, $E_k(t+1, w) = E_k(t, w)$. Let $U(w)$ be the set of all the rectangle indices j deleted while w is read. We are interested only when $U(w) \in \mathcal{U}_r$, which is not in general the case for an arbitrary word w satisfying only the three rules above. \square

We will now describe further rules that take account of coincident coordinates, and will prove in Lemma 5.4 that these rules are satisfied by $w = W(U)$, whenever $U \in \mathcal{U}_r$. It will turn out that, conversely, if w satisfies all the rules we present, then there is a unique $U \in \mathcal{U}_r$ such that $w = W(U)$.

Rule D: As rectangles are discarded, barrier values may change. A lefthand endpoint barrier value will not increase and may decrease. A righthand endpoint barrier value will not decrease and may increase. This rule will ensure that any rectangle whose removal would change the barrier value at i must be removed when the focus is no later than i . Let $j(t, k, w)$ be the \preccurlyeq_k -smallest element of $E_k(t, w)$, and let $\text{bv}(t, k, w) = r_{j(t, k, w), k}$. Suppose that $2 \leq i \leq 2d$ and that t_i is the first time at which the focus is at i . The condition on w is that, if $1 \leq k < i$, then $\text{bv}(t, k, w)$ is independent of t provided that $t_i \leq t \leq \text{length}(w)$. We denote this constant value by $\text{bv}(k, w)$.

Rule E: Fix i with $1 \leq i \leq 2d$. Let A be an endpoint list with ordering \preccurlyeq_i . By an A -block we mean a non-empty set of entries in A of the form

$$B = \{j | j \in A \text{ and } r_{ji} = x\}$$

for some $x \in \mathbb{R}$. Any $j \in A$ is contained in a unique A -block. Such a block can be of length one, but blocks can also be of any length up to N , as would happen, for example, if all of the d -rectangles in the list \mathcal{R} were equal to each other. We denote by $B_k(t, w)$ the initial $E_k(t, w)$ -block, and suppose t is such that the focus is at k . Let $b(t, w) = |B_k(t, w)|$. Then $b(t, w) > 0$. The condition on $w = w_t.v_t$ is that, if v_t starts with Δ , then $\delta(t, w) + b(t, w) \leq r$ and v_t starts with $\Delta^{b(t, w)}$. In other words, if v_t starts with a Δ , then all the elements of the initial block of $E_k(t, w)$ must be deleted before the focus increases.

Rule F: For each k with $1 \leq k \leq d$, $\text{bv}(2k - 1, w) \leq \text{bv}(2k, w)$. This rule ensures that only non-empty intersections are reached, and this can be checked as the focus leaves any even integer.

Lemma 5.4. *Let $U \in \mathcal{U}_r$ and $w = W(U)$. Then w satisfies all the rules above.*

Proof. From Equation 10 we immediately see that Rule A, Rule B and Rule C are true.

To prove Rule D, recall that $j(k, U)$ is the \preccurlyeq_k smallest element of $E_k(0, w) \setminus U$ (see Definition 5.3 for the meaning of $E_k(t, w)$). Since only elements of U are discarded while w is being read, $j(k, U)$ is never discarded. Therefore, for all t , $j(k, U)$ is the \preccurlyeq_k smallest element of $E_k(t, w) \setminus U$. Therefore, for all t ,

$$j(t, k, w) \preccurlyeq_k j(k, U). \quad (11)$$

We claim that, for $t \geq t_{k+1}$,

$$\text{bv}(t, k, w) = r_{j(t, k, w), k} = r_{j(k, U), k} = \text{bv}(k, U), \quad (12)$$

where the first and third equalities are definitions. For, if Equation 12 were false, then Equation 11 would give $j(t, k, w) \prec_k j(k, U)$. Recall from the definition just before Equation 8, that $U_k = \{j : j \prec_k j(k, U)\}$, so that $j(t, k, w) \in U_k$. However, for $t \geq t_{k+1}$, all elements of U_k have been discarded, in particular $j(t, k, w)$. This contradiction proves our claim, which proves Rule D.

Now we prove Rule E. Choose k with $1 \leq k \leq 2d$, and suppose the focus is at k at time t . Let B be the initial block of $E_k(t, w)$, and let $x \in \mathbb{R}$ be the real number such that $B = \{j \in E_k(t, w) | r_{jk} = x\}$. The initial Δ in v_t deletes some $j \in B \cap U^k$, so that $j \prec_k j(k, U)$. If k is even, this means that $x = r_{j, k} < r_{j(k, U), k}$, and, if k is odd, $x = r_{j, k} > r_{j(k, U), k}$. It follows that, for all $j \in B$, $j \prec_k j(k, U)$. But this means that $B \subseteq U_k$. Since the focus is on k , $B \cap U^j = \emptyset$ for $j < k$. We deduce that $B \subseteq U^k$, and so all elements of B are discarded immediately. This completes the proof of Rule E.

Rule F follows from Equation 12 and Lemma 4.3, and this completes the proof of Lemma 5.4. \square

Proposition 5.5. *The map $F_{WU} : \mathcal{U}_r \rightarrow \mathcal{W}_r$, defined by $F_{WU}(U) = W(U)$, is a bijection.*

Proof. Lemma 5.4 shows that we do indeed have $W(U) \in \mathcal{W}_r$, for each $U \in \mathcal{U}_r$.

Conversely, given $w \in \mathcal{W}_r$, we obtain a subset $U \subset \{1, \dots, N\}$ of discards by means of the process $\pi(w)$ (Definition 5.3). We need to prove the two conditions of Lemma 4.3.

Let $V_i = \{j : j \prec_i j(i, U)\}$. By the definition of $j(i, U)$ (see Definition 4.2), $V_i \subseteq U$, so, in due course, all elements of V_i are discarded during the process $\pi(w)$. It follows that, for i fixed and t sufficiently large (of course $t \leq r + 2d$ throughout), $j(t, i, w) = j(i, U)$, and therefore $\text{bv}(i, w) = \text{bv}(i, U)$. Rule F implies 4.3.1.

Let $j \in U$. Then, for some i , j is discarded while the focus is at i . While the focus is at i , only elements $k \in U$ with $k \prec_i j(i, U)$ are discarded. Therefore $j \prec_i j(i, U)$. We claim that $r_{j,i} \neq r_{j(i,U),i}$. For otherwise j and $j(i, U)$ would be in the same E_i -block. But then Rule E would imply that $j(i, U)$ is also discarded, implying that $j(i, U) \in U$. This contradiction implies that $j \prec_i j(i, U)$, which is 4.3.2. We have proved the two conditions of Lemma 4.3, and so $U \in \mathcal{U}_r$. \square

6 A brief sketch of the algorithm

We have reduced our problem to that of finding an algorithm that produces the words of \mathcal{W}_r one-by-one. Consider the finite set $\mathcal{W}(r)$ of words w satisfying Rule A, Rule B and Rule C. In the standard fashion, we form these words into a rooted tree $\mathcal{T}(r)$, with directed edges, each edge labelled with either Δ or S , so that there is exactly one path in \mathcal{T} from the root to a leaf for each word in $\mathcal{W}(r)$. There is a subtree \mathcal{T}_r corresponding to the words of \mathcal{W}_r .

Our algorithm traverses a certain subtree of the tree $\mathcal{T}(r)$, in depth-first order, choosing to follow a Δ -edge rather than an S -edge, where there is a choice. For efficiency reasons, we would prefer to arrange for this subtree to be exactly the subtree \mathcal{T}_r of interest. However, staying inside \mathcal{T}_r would require a substantial amount of additional computation—this is because Rule F cannot be easily checked term by term. Instead we explore a subtree of $\mathcal{T}(r)$ that is somewhat larger than \mathcal{T}_r . As soon as we discover that we are outside \mathcal{T}_r , we backtrack.

Our algorithm uses a recursive function call to perform a depth-first search of \mathcal{T}_r . We maintain a global data structure enabling us to check that the rules are being followed during the search. We denote the recursive call by $F(p, s)$. Here p is the maximum number of times we permit further rectangle indices to be discarded from $\{1, \dots, N\}$, that is, the maximum number of terms Δ that the algorithm may still generate, and p is mnemonic for *permissible discard*. Also s is the number of endpoint lists still to be examined, that is, the number of terms S still to be generated, and s is mnemonic for *shift*. The behaviour of $F(p, s)$ depends on the state of the global data structure, for example deciding whether to call $F(p - 1, s)$ or $F(p, s - 1)$ or both. The depth-first aspect of the search comes about since $F(p, s)$ calls $F(p - 1, s)$ before $F(p, s - 1)$ if the data structure tells us to call both. A principle followed by our algorithm is that $F(p, s)$ alters the most important part of the data structure before calling $F(p - 1, s)$ and/or $F(p, s - 1)$. On its return, $F(p, s)$ changes this part of the data structure back again before exiting. Thus, (the most important part of) the global data structure, when $F(p, s)$ starts, is the same as the structure when $F(p, s)$ finishes.

The depth-first search starts by calling $F(r, 2d)$.

7 Selection versus sorting

We will later describe in detail the data structures used by our algorithm to check, as rapidly as possible, that the rules are being followed. In this section we discuss a minor

improvement that can be made on the most obvious approach. The first three rules are enforced by the notation $F(p, s)$ and so have no implications for the data structure. The procedure described above requires us to order $\{1, \dots, N\}$ with respect to the $2d$ -different orderings \preceq_i described in Definition 3.5. This takes $O(dN \log N)$ steps.

However, we can do a little better than this. For each k with $1 \leq k \leq 2d$, let $G_k \subseteq \{1, \dots, N\}$ be the ordered set of the first $r+1$ elements of $\{1, \dots, N\}$ with respect to the ordering \preceq_k . It turns out (see Lemma 7.1) that our algorithm can be confined to consideration only of the ordered sets G_k , and that it will then run in the same number of steps (or slightly fewer). To compute the G_k we proceed as follows. For each k , find the $(r+1)$ -st element $p_k \in E_k$. An algorithm doing this is called a *selection algorithm*, and it requires $O(N)$ steps, as described in the next paragraph. As an unordered set $G_k = \{j : j \preceq_k p_k\}$ is found in $O(N)$ steps. G_k is then sorted in $O(r \log r)$ steps. To obtain all the G_k as ordered sets takes

$$O(d(N + r \log r)) \tag{13}$$

steps.

There is a randomized selection algorithm [2], due to Hoare, with an expected number of steps equal to $O(N)$, though the worst-case bound is $O(N^2)$. [1] contains an algorithm where the number of steps is bounded by $O(N)$ steps even in the worst case, though the constant hidden by the O -notation is not small. The latter algorithm only rarely outperforms the Hoare algorithm, and the Hoare algorithm also has the advantage that it can be carried out in place, that is, without needing more space than is needed by the input data. There is a good article about this topic, giving Hoare's code, at http://en.wikipedia.org/wiki/Selection_algorithm.

Given a word w in Δ and S satisfying Rule A, Rule B and Rule C, we have the process $\pi(w)$ (see Definition 5.3). We change this slightly to the process $\pi_G(w)$, which is the same as $\pi(w)$, except that, when deletion of j from E_k is required in $\pi(w)$, then no action is taken unless $j \in G_k$.

Lemma 7.1. *Rule D, Rule E and Rule F can be checked within the process $\pi_G(w)$.*

Proof. We fix k , with $1 \leq k \leq 2d$. By the time the process $\pi(w)$ has completed, α elements have been deleted from G_k , and β elements from $\{1, \dots, N\} \setminus G_k$, where $\alpha \leq \alpha + \beta \leq r$, so at most r elements have been deleted. Since $|G_k| = r + 1$, one or more elements of G_k is never deleted during $\pi(w)$. This means that we can compute $j(t, k, w)$ (defined in Rule D) within $\pi_G(w)$. So we can also compute $\text{bv}(t, k, w)$. We are therefore able to check Rule D and Rule F.

We assume now that w fails Rule E but passes all the other rules. We show that the failure can be seen within the process $\pi_G(w)$. Using the notation in which Rule E was stated, we suppose the rule fails at time t , when the focus is at k . We are assuming that $w = w_t.v_t$, where w_t has length t , and that v_t deletes the first element of $B_k(t, w)$, but not all of it. This means that v_t starts with $\Delta^\alpha.S$, where $1 \leq \alpha < b(t, w)$. The number of discards from G_k caused by w_t is bounded above by $\delta(t, w)$, and $\delta(t, w) + \alpha \leq r$, by Rule C. Therefore $w = w_t.v_t$ deletes no more than r elements from G_k , so that at least one element remains in G_k . Since $b(t, w) > \alpha$, the smallest such element must be in the block $B_k(t, w)$, so that we observe the failure of Rule E within G_k . \square

8 Data structure

We start with the given input: the $N \times 2d$ -matrix \mathcal{R} as in 2.4. Given r with $0 < r < N$, we want to generate all words in \mathcal{W}_r . In particular, this means that we need the ability to efficiently check the rules, avoiding unnecessary computation. We now describe the different parts of the data structure. As stated above, significant parts of the data structure, used and changed by $F(p, s)$, are restored as $F(p, s)$ exits. We will label these with the label **restore**.

8.1 **Tree**: The tree $\mathcal{T}(r)$ of all words satisfying our first three rules was introduced in Section 6. The subtree $\mathcal{T}_r \subseteq \mathcal{T}(r)$ is formed from \mathcal{W}_r , the set of words satisfying all six of our rules. **Tree** is a subtree of $\mathcal{T}(r)$ that normally grows inside \mathcal{T}_r as our algorithm progresses, and various words are explored. From time to time we may have to explore words outside \mathcal{W}_r , because it may not be immediately apparent that Rule F must inevitably fail, no matter how the word is extended. The error is eventually discovered, and the illegal branch of **Tree** is snipped off.

Each node ν of the tree corresponds to a unique word $w(\nu)$ in the terms Δ and S , read from left to right as we descend the tree. Each node of the tree is provided with data under the following headings:

8.2 **Nodes and pointers Δ , S , parent** : We establish two symbolic values **UNDEF** and **ILLEGAL**. The node ν has pointers $\nu.\Delta$, $\nu.S$ and $\nu.\text{parent}$, each pointing to another node of the tree or set equal to **UNDEF** or to **ILLEGAL**. As soon as we find out that $w(\nu)$ is not a proper prefix of a word in \mathcal{W}_r , we set $\nu.\Delta = \text{ILLEGAL}$, and $\nu.S = \text{ILLEGAL}$.

Initially, **Tree** has only one node, namely **root**, representing the null word, and initially all its three pointers point to **UNDEF**. For every node ν the **parent** pointer is never changed, so that the **parent** pointer of **root** always points to **UNDEF**. If $\nu \neq \text{root}$, then its **parent** pointer always points to a node, not to **UNDEF**, so this can be used to recognize **root**.

8.3 **CurrentNode (restore)**: is a global variable, pointing to the node ν of **Tree** currently reached by the algorithm. It fixes the word $w = w(\nu)$ corresponding to this node. Initially, **CurrentNode** points to **root**, the only initial node of **Tree**.

8.4 **BSZ (restore) and Q (constant)**: Let B_i be the set of G_i -blocks, arranged in order using \prec_i . We set $b_i = |B_i|$, and number the elements of B_i in order, using *block numbers* $\{1, \dots, b_i\}$, with 1 numbering the \prec_i -smallest block and b_i the \prec_i -largest block. Q is an $N \times 2d$ matrix, whose entry in position (j, i) gives the block number of the G_i -block containing the rectangle index j , provided $j \in G_i$, and is equal to **UNDEF** if $j \notin G_i$. Constructing Q takes $O(dN)$ steps and Q does not change during the algorithm. The changing block sizes are recorded in the $(r+1) \times 2d$ -matrix **BSZ**, such that $\text{BSZ}(k, i)$ is the block size of the k -th block of G_i if $k \leq b_i$. If $k > b_i$, then $\text{BSZ}(k, i) = \text{UNDEF}$. Initializing **BSZ** takes $O(dr)$ steps. Updating **BSZ** when the rectangle R_j is discarded or restored takes $O(d)$ steps.

8.5 **DLL (restore Doubly Linked List)**: In Section 7 we defined, for each i with $1 \leq i \leq 2d$, G_i as the \prec_i -ordered set of the $r+1$ rectangle indices that are smallest with respect to \prec_i . **DLL** is an $N \times 2d$ -matrix, whose k -th column is denoted by

DLL_k . Each entry of DLL is a pair of rectangle indices, that initially records G_i and its order as follows. If $j \notin G_i$, the entry at row j and column i is **UNDEF** and this entry is unchanged throughout the algorithm. If $j \in G_i$, the (j, i) entry is initially a pair $(\text{prev}, \text{next})$ of rectangle indices, specifying the ordering \preccurlyeq_i on G_i in the usual manner. We adjoin a row 0 to DLL to accommodate, for each i , a start entry at $(0, i)$ with $\text{prev} = \text{UNDEF}$ and next equal to the \preccurlyeq_i -smallest surviving element of G_i . If j is the \preccurlyeq_i -largest rectangle index surviving, then $\text{next} = \text{UNDEF}$ at position (j, i) . The entries in DLL are changed in the usual way (taking constant time) as elements $j \in G_i$ are deleted. When j is deleted from G_i , we do not change the temporarily meaningless values $(\text{prev}, \text{next})$ at position (j, i) ; instead these values are retained for use when j is restored. DLL takes $O(d(N + r \log r))$ steps to initialize. Deleting or restoring a rectangle index j takes $O(d)$ steps.

- 8.6 $\nu.\text{vol}$: Let m be the number of terms in the word $w = w(\nu)$ that are equal to S . For $k \leq m$, $\text{bv}(k, w) = r_{j,k}$, where j is the \preccurlyeq_k -smallest element remaining in the doubly linked list DLL_k —this is because Rule D is enforced at each step. For $2i \leq m$, let

$$\alpha(i) = \text{bv}(2i, w) - \text{bv}(2i - 1, w).$$

This is already known, unless $m = 2i$ and the final term of the word $w(\nu)$ is S . If $\alpha(i) < 0$, we set $\nu.\text{parent.S} = \text{ILLEGAL}$ and snip off the part of the tree hanging from ν . Otherwise set

$$\nu.\text{vol} = \prod \{ \alpha(i) \mid 2i \leq m \text{ and } \alpha(i) > 0 \},$$

which can be evaluated inductively, each step of the induction taking constant time.

- 8.7 $\nu.\text{dim}$: We set $\nu.\text{dim}$ equal to the number of i such that $\alpha(i) > 0$.
- 8.8 **Results**: For $0 \leq p \leq r$, **Results**(p) has three subcomponents, namely **dim**, **vol** and **address**. The first two subcomponents give the lexicographically largest result of the form (dim, vol) so far achieved by discarding p d -rectangles from the list \mathcal{R} , then taking the intersection of the remaining d -rectangles. The third subcomponent is the address of the **Tree** node ν_p that was current at the time that this best result was found. Initially each **dim** entry is -1 , each **vol** entry is **UNDEF**, and each **address** entry is equal to **UNDEF**. Initialization takes $O(r)$ steps.

9 Pseudocode

Here is pseudocode for the algorithm sketched in Section 6. We will use $\&$ and $*$ for the address and indirection operators, as in C. The address can be, as in C, an address in memory, or, in languages that do not attempt to mimic machine architecture, the number of a row in a table or matrix. The code starts with initialization, which consists of the sorting processes described above in Section 8 and setting up the various data structures in the obvious way. This takes $O(dN + dr \log r)$ steps.

The execution path of the program $F(p, s)$, whose description follows, depends on the *state* in which it starts. By the *state* we mean the values contained in the global data structures described above. So $F(p, s)$ stands for many different possible execution paths.

However, We will be able to give a reasonable upper bound for the number of steps that $F(p, s)$ requires, and this upper bound is independent of the particular execution path.

Rule A is enforced by the inequality $0 \leq s \leq 2d$. We need to check that $F(p, s)$ enforces the other five rules.

9.1 $F(p, s)$: deciding how the tree can be extended

Input: The global variables described above; p , an upper bound for the number of characters Δ still to be generated, or, equivalently, the number of rectangles still to discard; s , the number of characters S still to be generated, or, equivalently, $2d$ minus the number of characters S already generated; **CurrentNode**, the address of the current node ν in the tree.

Output: Assignment of the pointers $\nu.\Delta$ and $\nu.S$ to **ILLEGAL**, where appropriate. Storage of best results so far.

```

(1)       $f \leftarrow 2d - s + 1$ ; /* focus now at  $f$  */
(2)       $j \leftarrow$  smallest element of  $DLL_f$ ;
(3)       $\nu \leftarrow \text{CurrentNode}$ ; /*  $\nu$  is the current node */
(4)      if  $s == 0$ 
(5)          /* Exploration of this word is complete and all rules are
              * satisfied. */
              set  $\nu.\Delta = \nu.S = \text{ILLEGAL}$ ; /* Enforcing Rule B */
(6)          if  $\text{Results}(r - p).(\text{dim}, \text{vol}) < \nu.(\text{dim}, \text{vol})$  (lexicographically)
(7)               $\text{Results}(r - p).((\text{dim}, \text{vol}) \leftarrow (\nu.(\text{dim}, \text{vol}), \&\nu)$ ;
(8)          endif
(9)      else
(10)         /*  $s > 0$ . Check all rules, but without altering the main
              * data structures. */
              if  $p == 0$ 
(11)                  $\nu.\Delta \leftarrow \text{ILLEGAL}$ ; /* Enforcing Rule C */
(12)             else /*  $p > 0$  */
(13)                 if  $\text{BSZ}(Q(j, f), f) > p$ 
(14)                     /* Note that  $\text{BSZ}(Q(j, f), f)$  is the block containing
                        *  $j$ . When  $p = 0$ , the inequality is automatically sat-
                        * isfied, we have already discarded  $r$  rectangle indices,
                        * and further discards are not allowed. If  $p > 0$  and the
                        * inequality is satisfied then Rule E would be violated
                        * by a further discard and this rules out continuations
                        * of  $w(\nu)$  with next letter  $\Delta$ . */
(15)                      $\nu.\Delta \leftarrow \text{ILLEGAL}$ 
(16)                 else

```

```

/* Rule E is satisfied and we check Rule D */
(17)   foreach  $i$  such that  $1 \leq i < f$ 
(18)        $k \leftarrow$  smallest element of  $DLL_i$ ;
(19)       if  $Q(j, i) == Q(k, i)$  and  $BSZ(Q(j, i), i) == 1$ 
(20)           /*  $j, k$  in same  $G_i$ -block and Rule D would be
                * violated by a further deletion. */

(21)            $\nu.\Delta \leftarrow$  ILLEGAL;
(22)           break from foreach loop
(23)       endif
(24)   endfor
(25)   endif
(26)   endif
(27)   if  $f$  is even
(28)        $k \leftarrow$  smallest element of  $DLL_{f-1}$ 
(29)       if  $r_{k, f-1} > r_{j, f}$ 
(30)           /* Rule F would be violated by following with  $S$ . */
(31)            $\nu.S \leftarrow$  ILLEGAL;
(32)       endif
(33)   endif
(34) endif

```

9.2 $F(p, s)$: extending the tree, and recursion

At this stage we have determined all failures in the rules that would be immediately apparent on following $w(\nu)$ with either Δ or S . As a result, we know which new nodes to construct, and we proceed with this task. We also alter the global data structures as we explore the tree.

Input: The global variables; p , an upper bound for the number of characters Δ still to be generated, or, equivalently, the number of rectangles still to discard; s , the number of characters S still to be generated, or, equivalently, $2d$ minus the number of characters S already generated; **CurrentNode**, the address of the current node ν in the tree.

Output: Construction of new nodes of the tree. Determination of the best results possible so far, using recursion.

```

(1)    $f \leftarrow 2d - s + 1$ ; /* same as Line 1 of 9.1 */
(2)    $j \leftarrow$  smallest element of  $DLL_f$ ;

```

/* same as Line 2 of 9.1 */

```

(3)   $\nu \leftarrow \text{*CurrentNode}$ 
(4)  if  $\nu.\Delta \neq \text{ILLEGAL}$ 
(5)    Create node  $\alpha$ ;
(6)     $\text{Treelocation} \leftarrow \&\alpha$ ;
(7)     $\nu.\Delta \leftarrow \&\alpha$ ;
(8)     $\alpha.\text{parent} \leftarrow \&\nu$ ;
(9)    foreach  $i$  such that  $j \in \text{DLL}_i$ 
(10)      Delete  $j$  from  $\text{DLL}_i$ ;
(11)      Decrease  $\text{BSZ}(Q(j, i), i)$  by 1;
(12)    endfor
(13)     $F(p - 1, s)$ ; /*  $p > 0$  by Line 11 of 9.1 */
(14)    if  $\alpha.\Delta == \text{ILLEGAL}$  and  $\alpha.S == \text{ILLEGAL}$ 
(15)      set  $\nu.\Delta = \text{ILLEGAL}$ ;
(16)    endif
(17)    Reinsert  $j$  in the doubly linked lists  $\text{DLL}_i$ ;
(18)    Increase by 1 the various  $\text{BSZ}(Q(j, i), i)$  (see Line 11 above);
(19)     $\text{CurrentNode} \leftarrow \&\nu$ ;
(20)  endif
(21)  if  $\nu.S \neq \text{ILLEGAL}$ 
(22)    Create node  $\beta$ ;
(23)     $\text{Treelocation} \leftarrow \&\beta$ ;
(24)     $\nu.S \leftarrow \&\beta$ ;
(25)     $\beta.\text{parent} \leftarrow \&\nu$ ;
(26)    foreach  $i$  such that  $j \in \text{DLL}_i$ 
(27)      Delete  $j$  from  $\text{DLL}_i$ ;
(28)      Decrease  $\text{BSZ}(Q(j, i), i)$  by 1;
(29)    endfor
(30)    if  $f$  is even
(31)       $k \leftarrow$  smallest element of  $\text{DLL}_{f-1}$ ;
(32)      if  $r_{k, f-1} == r_{j, f}$ 
(33)         $\beta.(\text{dim}, \text{vol}) = \nu.(\text{dim}, \text{vol})$ ;
(34)      else /* By Line 31 of Subsection 9.1  $r_{k, f-1} < r_{j, f}$ . */
(35)         $\beta.\text{vol} \leftarrow \nu.\text{vol} \times (r_{j, f} - r_{k, f-1})$ ;
(36)         $\beta.\text{dim} \leftarrow \nu.\text{dim} + 1$ ;
(37)      endif
(38)    endif
(39)     $F(p, s - 1)$ ; /*  $s > 0$  by Line 5 of Subsection 9.1 */
(40)    if  $\beta.\Delta == \text{ILLEGAL}$  and  $\beta.S == \text{ILLEGAL}$ 
(41)      set  $\nu.S = \text{ILLEGAL}$ ;
(42)    endif
(43)    Reinsert  $j$  in the doubly linked lists  $\text{DLL}_i$ ;
(44)    Increase by 1 the various  $\text{BSZ}(Q(j, i), i)$  (see Line 11 above);
(45)     $\text{CurrentNode} \leftarrow \&\nu$ ;
(46)  endif

```


10 Estimate for number of steps

By induction on $p + s$, there is a unique function $f(p, s)$ with the following properties:

$$\begin{aligned} f(p, 0) &= 1 \text{ for } p \geq 0; \\ f(0, s) &= s + 1 \text{ for } s \geq 0; \\ f(p, s) &= f(p - 1, s) + f(p, s - 1) + d \text{ for } p > 0 \text{ and } s > 0. \end{aligned} \tag{14}$$

From the pseudocode in Section 9, one sees that there is a constant $k > 0$, independent of N , d and r , such that the number of steps necessary to execute $F(p, s)$ is bounded by $k \cdot f(p, s)$.

Using the addition formula for binomial coefficients

$$\binom{n}{s} = \binom{n-1}{s} + \binom{n-1}{s-1},$$

one checks that $f(p, s)$ is given by

$$f(p, s) = d \binom{p+s}{s} + \binom{p+s+1}{s} - d, \tag{15}$$

since it satisfies Equation 14. In particular, we are interested in

$$f(r, 2d) = d \binom{r+2d}{2d} + \binom{r+2d+1}{2d} - d. \tag{16}$$

To complete the computation of the upper bound for the number of steps involved in our algorithm, we need to extract the information in the data structure **Results** (see 8.8). From $\rho = \mathbf{Results}(p)$, we find the set of discarded d -rectangles by following **parent** pointers, starting at $\rho.\mathbf{address}$. This takes at most $O(p + 2d)$ steps. Summing over $0 \leq p \leq r$, we obtain a bound

$$O((r+1)2d + r(r+1)/2). \tag{17}$$

steps, where the constant implicit in the O notation is independent of N , d and r . Combining this with the bounds of Equations 13 and 16, we get an overall bound

$$O\left(d \binom{r+2d}{2d} + dN\right) \tag{18}$$

since the omitted terms are dominated by those that remain in Equation 18.

11 The dual situation

Let \mathcal{S} be the set of closed non-empty d -rectangles in \mathbb{R}^d , together with the null-set. We have a lattice: the greatest lower bound of a finite subset of \mathcal{S} is given by intersection. The least upper bound of any finite subset of \mathcal{S} exists, but is not in general equal to the union.

Theorem 2.5 applies equally well to the dual situation. Here we are given N d -rectangles and a number r with $0 \leq r \leq N$. The task is to find which r rectangles to discard, so that the smallest rectangle containing all except these r rectangles is as small as possible.

We can then prove the dual result to Theorem 2.5 by reversing all inequalities and replacing intersection by least upper bound. It is possible to sketch scenarios in which such a result might be used, but we spare the reader rather than labour the point.

References

- [1] M. Blum, R.W. Floyd, V. Pratt, R.L. Rivest, and R.E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7(4):448–461, 1973.
- [2] C.A.R. Hoare. Algorithm 63: partition algorithm 65: Find. *Communications of the ACM*, 4(7):321–322, 1961.
- [3] Shan-e-Ahmed E. Raza, Ahmad Humayun, Sylvie Abouna, Tim W. Nattkemper, David B.A. Epstein, Michael Khan, and Nasir M. Rajpoot. Ramtab: Robust alignment of multi-tag bioimages. *Plos One*, 7(2):e30894–e30894, 2012.
- [4] Walter Schubert, Bernd Bonnekoh, Ansgar J. Pommer, Lars Philipsen, Raik Bckelmann, Yanina Malykh, Harald Gollnick, Manuela Friedenberger, Marcus Bode, and Andreas W. M. Dress. Analyzing proteome topology and function by automated multidimensional fluorescence microscopy. *Nature biotechnology*, 24(10):1270–8, 2006.